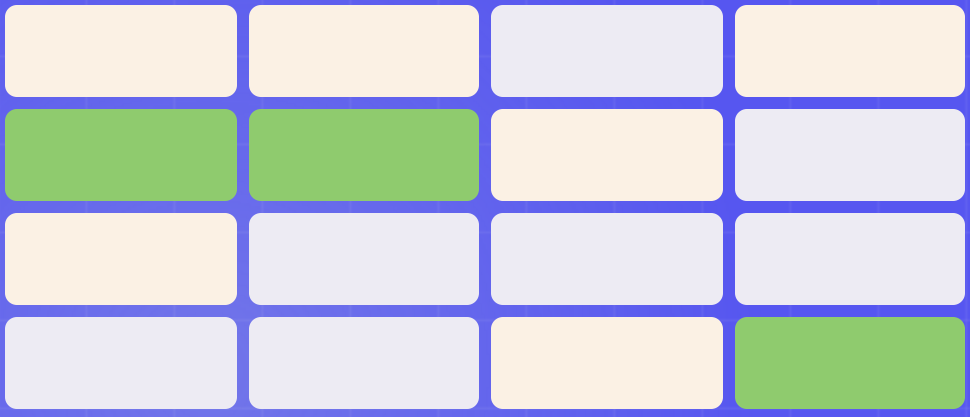
VelocityEngine

EBOOK

Your competitive coverage map is missing the conversations that matter

Which buyer conversations is a named competitor running unopposed, by persona and funnel stage?



PERSONA × STAGE

FOR PRODUCT MARKETING LEADERS

TL;DR

01 Your content reports show that an asset exists for a buyer. They don't show whether it shapes how that buyer sees the problem. That gap between presence and occupancy is where competitive deals are lost.

02 A persona-by-stage coverage map shows which buyer conversations you occupy, which you're only present in, and which are empty. An asset inventory can't produce that finding.

03 Read a competitor the same way. Topic overlap tells you little. When a competitor opens asset after asset for one persona with the same tension, they're claiming that conversation on purpose.

04 Your own gaps aren't a talent problem. When positioning lives in a reference document, every asset starts from someone's interpretation of it, and the framing drifts.

05 The fix is structural: make positioning the source each asset is built from. Paysafe reported 3–5x faster campaign delivery, 75% less review time and 3–4x more persona-based content after making that shift.

Table of contents

SYNOPSIS

This ebook shows product marketing leaders how to find the buyer conversations a competitor is running unopposed: map your own coverage by persona and stage, read competitor content for repeated tension framing, trace why your framing drifts, and close the gap by building assets from governed positioning. It ends with a 30-minute self-test you can run with your team.

— Introduction 1

01 Expose the coverage illusion your metrics can't see 2

02 Map buyer coverage at the persona-stage level 5

03 Turn competitor content into a coverage intelligence signal 8

04 Trace the drift before you blame the brief 12

05 Stop consulting your positioning. Inherit it. 15

06 Conclusion 17

✓ The 30-minute self-test 19

Introduction

Most Directors of Product Marketing can describe their competitive position in detail. They can name the positioning pillars, cite the coverage framework and pull up the asset library. The confidence is real. But ask one question (which buyer conversations is a named competitor running unopposed, by persona and by funnel stage?) and that confidence changes shape. Not because the answer is complicated, but because the reports they rely on were never built to produce it.

The standard fix is to audit keyword parity and map asset volume to funnel stages. Teams that run that audit usually find no missing topics. What they can't find is which persona-stage conversations a competitor has claimed by framing the buyer's problem the same way, across formats, until buyers associate that framing with that company. A sales debrief surfaces the loss. The coverage report never did.

THIS EBOOK GIVES YOU A WAY TO SEE IT. YOU'LL LEARN TO:

- tell the difference between having content for a buyer and owning that buyer's conversation;
- map your own coverage by persona and stage, and mark where you're present but not occupying;
- read a competitor's content for intentional narrative ownership, not just topic overlap;
- trace why your own framing drifts, and test whether your positioning structure can close the gap.

At the end, a 30-minute self-test puts all three diagnostics on one page so you can run them with your team.

01

Expose the coverage illusion your metrics can't see

Why keyword and traffic reports overstate competitive coverage

According to VelocityEngine's content coverage benchmark, most teams cover only 10 to 30% of the segments and personas sales is pursuing, and they can't see where they're exposed. That isn't a content volume problem. It's a measurement problem. The dashboards look complete, and the gap isn't visible from inside them. This is the coverage illusion: a measurement system that confirms presence while hiding occupancy.

THREE TERMS USED IN THIS EBOOK

Presence: a campaign asset exists for a topic, persona or funnel stage.

Occupancy: your narrative has shaped how a buyer frames their problem before they evaluate solutions.

Tension framing: the sentence or paragraph that names why a specific buyer has a problem. Repeating the same tension for the same persona at the same stage is how a company builds occupancy.

A keyword audit can confirm presence. It can't confirm occupancy. The two get confused when a visible metric stands in for an invisible one and nobody checks whether the substitution holds.

Here's how that substitution works. Intent signals and keyword rankings measure demand that already exists. They confirm that buyers searched for a category. They don't tell you whether your positioning shaped how those buyers framed their problem. A buyer who found your content and a buyer whose problem was framed by your narrative are not the same buyer, but your reporting treats them the same.

So keyword parity doesn't equal narrative ownership. Teams miss this because they test for topic coverage while the contest runs at the persona-stage level.

TRY THIS WITH YOUR TEAM

Ask a product marketing manager to name, without opening a dashboard, which funnel stage your messaging occupies for your second-tier persona in your top competitor's strongest vertical. Don't accept "we have content for that." The tell isn't a wrong answer. It's the pivot: reaching for asset counts, keyword rank or traffic share instead of an answer. That pivot means your measurement has no way to show the thing you asked about.

Fixing this has a cost, and it's worth naming. Rebuilding coverage measurement around persona-stage occupancy breaks continuity with existing reports. Historical benchmarks stop being comparable, and board dashboards built on keyword and traffic signals don't translate cleanly to a model organized by persona, stage and competitive narrative. That friction is the price of trading a measurement system that feels complete for one that is.

The illusion shows most clearly after a lost deal. The team runs a keyword audit, finds parity (same topics, comparable traffic) and concludes the problem is content quality or volume. The audit had no way to show that the competitor already owned the buyer's framing at the evaluation stage. The assets were present. The conversation wasn't.

Presence and occupancy aren't degrees of the same thing. They're measured differently, and they signal different things.

The next step is to measure the one your reports leave out.

02

Map buyer coverage at the persona-stage level

Step one of a competitor content analysis framework: your own map

> **What our sales leaders are struggling with is: do we have the right assets, and are we delivering the right point of view to the right personas? That's where we're struggling. I've been looking for something to give proof.**
>
> BRYANT BELL, PRODUCT MARKETING MANAGER, EXTERRO

That question (not "do we have content?" but "does our content occupy the right conversation?") is what a persona-stage coverage map answers. Here's how to build one for a single ICP.

1 Set the grid

Rows are the personas in the buying committee for your primary ICP: the ones sales actually meets, not every title in your persona document. Columns are the stages of the buying journey: awareness, consideration, selection and commit. Keep the first pass to four to six personas.

2 Write the target tension for each cell

For every persona-stage cell, write one sentence that names the pressure that buyer feels at that moment, in their words. A CFO at consideration doesn't have "visibility challenges." They have a board meeting where they can't explain why forecast and actuals diverged.

If you can't write the sentence for a cell, mark it. That's already a finding.

3 Place your assets

Put each live asset in the cell it was built for. Don't just count them. Read the opening of each one and note the tension it leads with.

4 Score each cell

Use three states:

Occupied	**Occupied:** assets in the cell open with the target tension, consistently, in language close to the sentence you wrote.
Present	**Present, not occupying:** assets exist, but they open with different framings, or with a generic tension shared by the whole buying committee.
Empty	**Empty:** nothing is built for this buyer at this stage.

5 Rank the gaps

Mark the cells that matter most to pipeline: the personas sales loses deals with, and the stages where deals stall. Those are the cells the next two chapters focus on.

	AWARENESS	CONSIDERATION	SELECTION	COMMIT
CFO	Present	Present	Empty	Present
VP of Operations	Occupied	Occupied	Present	Empty
Head of Compliance	Present	Empty	Empty	Empty
Procurement lead	Empty	Empty	Present	Occupied

An illustrative persona-stage coverage map for one ICP. “Present” means present, not occupying.

What the map usually shows

The biggest risk is rarely the empty cells. Teams usually know where they’ve built nothing. The risk is the cells marked “present, not occupying,” where the asset count says you’re covered and the buyer hears four different stories. In the example above, the CFO looks covered at three of four stages. None of those cells is occupied. Those are the conversations a competitor with one consistent narrative can take without outspending you.

You can build a map like this by hand for one ICP in an afternoon. Keeping it current across every segment is where most teams stall.

03

Turn competitor content into a coverage intelligence signal

Competitor content analysis for competitive positioning

Most teams don't lack competitive intelligence. The intelligence they collect answers the wrong question. Keyword overlap tells you where you share traffic. Topic frequency tells you what a competitor writes about. Neither tells you which buyer conversations they've claimed, with which personas, at which stage, consistently enough that buyers now associate that framing with that company.

Read a competitor's content at the topic or keyword level and the signal is thematic. They cover compliance. They cover workflow efficiency. They cover total cost of ownership. At that level, most competitors in a mature category look similar. The thematic read confirms overlap. It doesn't reveal occupancy.

Read the same content through the persona-stage lens from chapter 2 and the signal changes. The question shifts from "what do they write about?" to "which role-specific tension do they open with, and have they opened with it before?" Repetition of the same role-specific tension is the behavioral signal of intentional coverage.

A competitor who opens three consecutive consideration-stage assets for the same persona with the same pressure point isn't being repetitive. They're building ownership of that conversation, and no tool that groups content by topic, keyword or intent will show it.

> **Every content decision is informed by intuition or dated information. We just throw content at the wall. This illuminates gaps and what competitors are doing, and it radically accelerates reaction time from marketing.**
>
> CHRIS HUTCHINS, HEAD OF CONTENT, AUDIENSE

"Throw content at the wall" isn't self-deprecation. It describes what broad competitive monitoring produces: a reaction posture, not a coverage posture.

This read takes effort. A structured persona-stage read of a competitor's content needs PMM time and a disciplined framework, while broad monitoring of keywords, messaging themes and launch cadence is faster and easier to maintain. The cost of choosing speed is that you can track everything a competitor publishes and never notice they've built a consistent narrative for your highest-value persona at the stage where your win rate matters most.

A HYPOTHETICAL EXAMPLE

A PMM reads a competitor's last two quarters of content at the topic level and finds broad similarity to their own. Both companies cover the same buyer tensions. At the persona-stage level, the finding changes: the competitor opened every major consideration-stage asset for one persona with the same role-specific tension. The PMM's own content for that persona at that stage opened with four different framings over the same period. The competitor isn't outspending them. They're out-covering them on a single buyer conversation, on purpose, repeatedly.

How to run the read

1. Pick one competitor: the one that creates the most friction in your deals.
2. Choose one cell from your map, ideally one you marked "present, not occupying" or empty, and pull the competitor's last two quarters of content for that persona and stage.
3. For each asset, write down the tension it opens with.
4. Test whether that tension is generic to the buying committee or specific to the persona's role.
5. If it's role-specific, check whether earlier assets for that persona and stage opened with the same tension.
6. Repetition of a role-specific tension across several assets is the confirmation signal. Topic consistency is not.

What this read does and doesn't prove

It doesn't prove the competitor is winning because of this pattern. The evidence shows intentional persona-stage occupancy, not whether that occupancy converts to pipeline or closes deals. Win-rate data tells you whether the conversation is working. You need both reads to draw a causal conclusion.

What it does prove is that your current monitoring can't see the pattern at all. You can track every asset a competitor publishes and still miss that they've claimed a conversation you haven't entered, not because they outspent you, but because they repeated themselves with precision while your team varied its framing by instinct.

The coverage gap isn't always where you have no assets. Sometimes it's where you have assets with no consistent persona-specific narrative.

Which raises the obvious question: if a competitor can hold one framing steady, why can't your team?

04

Trace the drift before you blame the brief

Why messaging drifts between strategy and execution

The last two chapters showed where the gaps are. This one explains how they form: not in the strategy room, but in the space between strategy and execution.

The positioning document didn't fail. The campaign team didn't ignore it. The drift happened because the two never shared the same moment of creation.

When positioning lives in a document that campaign teams reference at briefing time but don't build from at creation time, every asset begins with an approximation. The writer reads the brief. The campaign manager interprets the brief. The regional team adapts both. Each compression is locally reasonable. Together, they wear the message down. By the time the asset ships, the persona-specific tension the PMM defined may be present in spirit, but absent in the precise language that makes a CFO, a head of compliance or a VP of operations recognize their own problem.

This isn't a talent problem. It's a distribution problem.

The mechanism runs in one direction. Positioning gets written once, stored centrally and referenced now and then. Creation happens continuously, in distributed contexts, under deadline pressure. The gap between those two rhythms is where drift builds up. Writers don't paraphrase out of carelessness. They paraphrase because a document is a reference, not a governed source. A reference gets interpreted. A governed source gets inherited.

The consequence isn't just inconsistent messaging. It's a competitive positioning failure.

What drift looks like at the asset level

Pull three shipped assets from three different campaigns that target the same persona at the same stage. Read the tension statement in each. Check whether each traces back to the same root positioning language. The sign of drift isn't contradiction. It's variation that's directionally similar but not identical: each asset sounds like the brand, but no two frame the buyer's tension the same way.

A CONSTRUCTED EXAMPLE

A PMM defines a specific tension for a CFO at the consideration stage. Three campaigns ship that quarter targeting that persona. One asset frames the tension as a reporting-lag problem. Another frames it as workflow fragmentation. A third frames it as a data-trust problem. All three fit the positioning. None of them is the tension the PMM defined. A competitor running one consistent narrative for that persona at that stage has, by default, become the company that frames the conversation, because the three framings split the signal and none gets the repetition it needs to shape how buyers see the problem.

That competitor didn't outwork the team. They out-structured it.

Decentralized creation has real benefits. Local teams adapt messaging, ship faster and coordinate less, and that speed is real. But every locally reasonable adaptation is a vote against positioning coherence, and the PMM review cycle that follows (rewriting assets, correcting tone, chasing consistency) isn't quality control. It's the debt of a system that never governed the source.

That debt comes due when the market moves faster than the review cycle can correct it. A competitor launches a new narrative aimed at your persona at your stage. Your team responds with three assets that each frame the competitive tension differently. None of them lands, not because any one asset is weak, but because the signal is fractured before it reaches the buyer.

Drift doesn't show up in campaign dashboards, click-through rates, conversion rates or pipeline attribution. It shows up only when you hold three assets side by side and ask: do these all frame the same buyer tension, in the same words, with the precision the PMM intended? Most teams don't run that test. They run the asset, and the drift compounds campaign by campaign.

05

Stop consulting your positioning. Inherit it.

Operationalizing competitive positioning at the point of creation

Paysafe reported 3–5x faster campaign delivery, 75% less review time and 3–4x more persona-based content after changing where positioning lives when content is created.

3–5x
faster campaign delivery

75%
less review time

3–4x
more persona-based content

That didn't come from hiring more PMMs. It's a structural result, not a productivity one.

The earlier chapters showed how to find the persona-stage gaps a competitor is exploiting and why your own framing drifts. The question now is what happens the moment your team sits down to close a gap.

Under a briefing-and-reference model, the PMM defines the tension. The brief summarizes it. The writer interprets the summary. The asset reflects the interpretation. By the time the buyer reads it, the specific framing the PMM built has been approximated twice. The drift isn't obvious in the asset, because the direction is usually right. But the precise tension is gone.

A governed foundation doesn't change how positioning is built. It changes where positioning lives while content is created. When every campaign asset, including AI-generated content, draws from the same foundation, the persona-specific tension, the stage narrative and the competitive framing don't have to be re-approximated by whoever writes the asset. AI-generated content is only useful when it keeps that nuance and stays reviewable, governed and aligned to the target segment. Built this way, the foundation closes the gap between what the PMM defined and what the buyer reads.

IN VELOCITYENGINE

That foundation is the Strategy Foundation. Campaign briefs auto-populate from it, and content is created against it, so a gap you find on the map can be closed with assets that carry the framing you defined.

Hutchins noted that while the gap analysis alone is valuable, being able to generate a targeted campaign for a specific gap in under a minute changes how fast the team can operate. That speed isn't a workflow tweak. It comes from removing the interpretive step.

06

Conclusion

Three steps to run before your next launch cycle

1 Map your own coverage

Build the persona-stage map for your primary ICP (chapter 2). The goal isn't to inventory what exists. It's to find every cell where content is present but doesn't name the buyer's role-specific tension. If the content in a cell doesn't name the pressure a VP of operations or a procurement lead feels at the moment of evaluation, it doesn't occupy their conversation. Mark those cells as space a competitor may already own.

2 Read the competitor who causes the most friction

Run the structured read from chapter 3 on those cells. Look for tension-framing consistency, not topic overlap or keyword density. Repetition of one buyer tension across several assets, for one role at one stage, is the sign of intentional narrative ownership. That's the gap that closes deals for them, not content volume.

3 Test your positioning architecture

Check whether you can trace a direct line from your positioning to a shipped asset without an interpretive step (chapter 4). If a campaign team needs a brief to translate your positioning, the gap is structural. A governed foundation makes positioning the source an asset is built from, not a document it's checked against. Find where that chain breaks and define the change before your next launch cycle starts.

Teams that run this sequence close the gap before it compounds. Teams that don't will keep producing well-crafted assets that lose deals to competitors who own the conversation.

The 30-minute self-test

Run these three tests with your team. Ten minutes each.

Test 1: The dashboard pivot (chapter 1)

10 MIN

Ask a PMM, without a dashboard: which stage does our messaging occupy for our second-tier persona in our top competitor's strongest vertical?

- [ ] They named a stage and the tension we own there
- [ ] They reached for asset counts, keyword rank or traffic

Test 2: The competitor read (chapter 3)

10 MIN

Persona and stage:

Tension the competitor opens with (last three assets):

- [ ] The same role-specific tension, repeated
- [ ] Generic or varied framing

Test 3: The three-asset drift check (chapter 4)

10 MIN

Tension each of our three assets opens with, for the same persona and stage:

- [ ] The same tension, in the same words
- [ ] Directionally similar, but not the same

Reading your results: a pivot in test 1 means your reporting can't see occupancy. Repetition in test 2 means the competitor is claiming that conversation on purpose. Variation in test 3 means your positioning is being consulted, not inherited.

You've seen how to find the buyer conversations a competitor is running unopposed. By hand, that works for one ICP.

VelocityEngine maps your content coverage by segment, persona and funnel stage, runs the same analysis on competitor content, and carries your positioning into every campaign through a governed Strategy Foundation.

See where you stand. Get your free Content Coverage Score.

Get your coverage score ›

velocityengine.co/content-coverage-heatmap-signup

VelocityEngine

© 2026 VelocityEngine. All rights reserved.